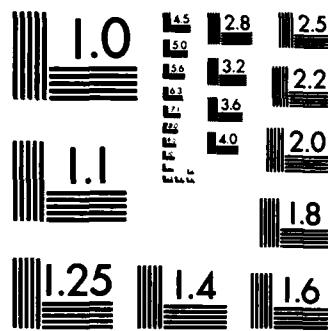


HD-A134 082 PROGRAM ABSTRACTS SPARSE MATRIX SYMPOSIUM 1982 HELD AT 1/1
FAIRFIELD GLADE TENNESSEE OCTOBER 24-27 1982(U) OAK
RIDGE NATIONAL LAB TN Y A JACOBS 1982

UNCLASSIFIED

F/G 12/1 NL

END



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A



AD-A134082

PROGRAM and ABSTRACTS

3

SPARSE MATRIX SYMPOSIUM 1982

Fairfield Glade, Tennessee
October 24-27, 1982

DTIC
ELECTED
S OCT 17 1983 D

DTIC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

SPONSORS

- U.S. Army Research Office
- U.S. Office of Naval Research
- Oak Ridge National Laboratory

(Mathematics and Statistics Research Department)

83 10 12 006

SYMPOSIUM CHAIRMAN

Robert C. Ward
Union Carbide Corporation
Oak Ridge National Laboratory

ORGANIZING COMMITTEE

Robert E. Funderlic
Union Carbide Corporation
Oak Ridge National Laboratory

Robert J. Plemmons
North Carolina State
University

Michael T. Heath
Union Carbide Corporation
Oak Ridge National Laboratory

David S. Scott
University of Texas - Austin

ADVISORY COMMITTEE

Jane K. Cullum
IBM T. J. Watson Research
Center

William G. Poole, Jr.
Boeing Computer Services

Carlos A. Felippa
Lockheed Palo Alto Research
Laboratory

Donald J. Rose
Bell Laboratories

J. Alan George
University of Waterloo

Martin H. Schultz
Yale University

Gene H. Golub
Stanford University

Margaret H. Wright
Stanford University

CONFERENCE COORDINATOR

Norma F. Callaham
Oak Ridge National Laboratory

TECHNICAL EDITOR

Vivian A. Jacobs
Oak Ridge National Laboratory

SECRETARIAL ASSISTANCE

Mitzy Denson
Union Carbide Corporation
Oak Ridge National Laboratory

Carole Yount
Union Carbide Corporation
Oak Ridge National Laboratory

Accession For	
NTIS GRAAI	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By Per DTIC Form 50	
Distribution/on file	
Availability Codes	
Avail and/or	
Dist	Special
A	



PROGRAM AND ABSTRACTS

Sparse Matrix Symposium 1982

Fairfield Glade, Tennessee

October 24-27, 1982

Robert C. Ward, Chairman

SPONSORS

U.S. Army Research Office

U.S. Office of Naval Research

Oak Ridge National Laboratory

(Mathematics and Statistics Research Department)

Prepared by Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830
 Operated by Union Carbide Corporation - Nuclear Division for the U.S. Department of Energy

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

TABLE OF CONTENTS

Final Program	1
Timetables	11
Author Index	17
Invited Abstracts	23
Contributed Abstracts	33

FINAL PROGRAM

SPARSE MATRIX SYMPOSIUM
Fairfield Glade, Tennessee
October 24-27, 1982

Sunday, October 24

6:00-10:00 p.m. Registration.

8:00-10:00 p.m. Reception.

Monday, October 25**

8:45 a.m. Welcoming Remarks, Robert C. Ward.

Session 1. Invited Papers. St. George Room
Chairperson: Robert C. Ward

9:00 a.m. Iain S. Duff* and John K. Reid, "Direct Methods for
Solving Sparse Systems of Linear Equations."

9:45 Stanley C. Eisenstat, "Iterative Methods for Solving Large
Sparse Linear Systems."

10:30 Coffee break.

Session 2A. Contributed Papers. Druid-Dorchester Room
Chairperson: J. Alan George

11:00 a.m. Albert M. Erisman, Roger G. Grimes, John G. Lewis*, and
William G. Poole, Jr., "A Structurally Stable Modifi-
cation of the Hellerman-Rancic Algorithm for Reordering
Unsymmetric Sparse Matrices."

11:20 Arne Drud, "Spike Selection for Large Sparse Sets of
Nonlinear Equations."

11:40 Robert Fourer, "Solving Sparse Staircase Systems by
Gaussian Elimination."

*Presenter of paper.

**Poster Session open during all sessions and breaks.

Monday, October 25**

Session 2B. Contributed Papers. Druid-Dorchester Room
Chairperson: Carlos A. Felippa

11:00 a.m. R. L. Cox, "LSODE28: A Variant of LSODE for Problems
Having a General Large Sparse Jacobian."

11:20 Granville Sewell, "IMSL Software for Differential
Equations in One Space Variable."

11:40 J. W. Neuberger and R. J. Renka*, "A Portable Software
Package for Nonlinear Partial Differential Equations."

12:00 Break.

Session 3A. Contributed Papers. St. George Room
Chairperson: J. Alan George

12:10 p.m. Robert Koury, David F. McAllister, and William J.
Stewart*, "Block Iterative Methods with Aggregation for
Solving Nearly Completely Decomposable Markov Chains."

12:30 P. M. Dearing, D. R. Shier, and D. D. Warner*, "Maximal
Chordal Subgraphs and Derived Matrix Splittings."

12:50 Linda Kaufman, "Fast Direct Methods for Other Sparse
Matrix Problems."

Session 3B. Contributed Papers. Druid-Dorchester Room
Chairperson: Carlos A. Felippa

12:10 p.m. Ken Kaneko, "The Turn-Back LU Procedure for Computing a
Sparse and Banded Basis of the Null Space."

12:30 M. T. Heath, R. J. Plemmons*, and R. C. Ward, "Sparse
Orthogonal Schemes for Structural Optimization Using the
Force Method."

12:50 A. S. Rao, "Sparse Matrices in Simplification of Systems."

1:10 Lunch, informal afternoon sessions, recreation.

Monday, October 25**

Session 4A. Contributed Papers. St. George Room

Chairperson: John K. Reid

7:30 p.m. Tom Manteuffel* and Vance Faber, "Conjugate Gradient Implies Normal."

7:50 Anne Greenbaum, "Analysis of a Multigrid Method as an Iterative Technique for Solving Linear Systems."

8:10 A. Brandt, S. McCormick*, and J. Ruge, "Algebraic Multigrid (AMG) for Geodetic Equations."

Session 4B. Contributed Papers. Druid-Dorchester Room

Chairperson: Jack J. Dongarra

7:30 p.m. D. A. Calahan, "Sparse Direct Methods for Vector Multiprocessors."

7:50 Ronald D. Coleman* and Edward J. Kushner, "The Sparse Matrix Library for the FPS-164 Attached Processor."

8:10 A. Sameh* and C. Taft, "Preconditioning Strategies for the Conjugate Gradient Algorithm on Multiprocessors."

8:30 Break.

Session 5. Invited Paper. St. George Room

Chairperson: David S. Scott

8:45 p.m. B. N. Parlett, "Software for Sparse Eigenvalue Problems."

9:30 End of evening session.

Tuesday, October 26**

Session 6. Invited Papers. St. George Room

Chairperson: Robert J. Plemmons

8:45 a.m. Michael T. Heath, "Numerical Methods for Large Sparse Linear Least Squares Problems."

9:30 Allen J. Pope, "Geodetic Computations and Sparsity."

10:15 Coffee break.

Session 7A. Contributed Papers. St. George Room

Chairperson: William G. Poole, Jr.

10:40 a.m. Aram K. Kevorkian*, Fred G. Gustavson, and Gary D. Hachtel, "A New Theory on Permuting Matrices to Block Triangular Form."

11:00 Fred G. Gustavson*, Gary D. Hachtel, and Aram K. Kevorkian, "An Improved Assignment Algorithm."

11:20 Gary D. Hachtel*, Fred G. Gustavson, and Aram K. Kevorkian, "A 1-Pass Procedure for Maximal Assignment and Finding the BLT Form."

11:40 Fred G. Gustavson, "A New Version of Tarjan's Strong Connect Algorithm."

Session 7B. Contributed Papers. Druid-Dorchester Room

Chairperson: Gene H. Golub

10:40 a.m. David S. Scott, "LAS02-Sparse Symmetric Eigenvalue Package."

11:00 Jane Cullum* and Ralph A. Willoughby, "An Accelerated Block Lanczos Procedure for Extreme Eigenvalues of Symmetric Matrices."

11:20 Paul S. Jensen, "A Production Eigenanalysis System for Large, Generalized Symmetric Problems."

11:40 Horst D. Simon, "The Lanczos Algorithm with Partial Reorthogonalization for the Solution of Nonsymmetric Linear Systems."

12:00 Break.

Tuesday, October 26**

Session 8A. Contributed Papers. St. George Room

Chairperson: William G. Poole, Jr.

12:10 p.m. Silvio Ursic, "Inverse Matrix Representation with One Triangular Array (Implicit Gauss)."

12:30 John de Pillis* and Wilhelm Niethammer, "A Stationary Iterative Method that Works Even for Hermitian $Ax = b$."

12:50 C. R. Johnson, "The Classes $M_{n,k}$ and Properties of Sparse Matrices Resembling Those of Matrices in a Smaller Dimension."

Session 8B. Contributed Papers. Druid-Dorchester Room

Chairperson: Gene H. Golub

12:10 p.m. Alan George and Joseph Liu*, "Row Ordering Schemes for Sparse Givens Transformations."

12:30 Alan George and Esmond Ng*, "Solution of Sparse Under-determined Systems of Linear Equations."

12:50 M. S. Kamel* and K. Singhal, "An Efficient Algorithm for the Factorization of Possibly Rank Deficient Matrices."

1:10 Lunch, informal afternoon sessions, recreation.

Session 9A. Contributed Papers. St. George Room

Chairperson: Linda Kaufman

7:30 p.m. Iain S. Duff and John K. Reid*, "The Multifrontal Solution of Unsymmetric Sets of Linear Equations."

7:50 John R. Gilbert* and Robert Schreiber, "Nested Dissection with Partial Pivoting."

8:10 Albert M. Erisman, "Matrix Modification and Partitioning."

Tuesday, October 26**

Session 9B. Contributed Papers. Druid-Dorchester Room

Chairperson: Yueh-er Kuo

7:30 p.m. A. Behie and P. A. Forsyth Jr.*, "Incomplete Factorization Methods for Fully Implicit Simulation of Enhanced Oil Recovery."

7:50 W. I. Bertiger*, D. A. Calahan, P. T. Woo, "The Effect of Computer Architecture on Direct Sparse Matrix Routines in Petroleum Reservoir Simulation."

8:10 W. P. Kamp, "A Three Dimensional Finite Element Magneto-telluric Modeling Program."

8:30 Break.

Session 10. Invited Paper. St. George Room

Chairperson: Jane K. Cullum

8:45 p.m. W. M. Coughran, Jr., W. Fichtner, E. H. Grosse, and D. J. Rose*, "Numerical Simulation of VLSI Circuits."

9:30 End of Evening Session.

Wednesday, October 27

Session 11. Invited Papers. St. George Room

Chairperson: Robert E. Funderlic

8:45 a.m. Dianne P. O'Leary, "Solving Mesh Problems on Parallel Processors."

9:30 Philip E. Gill, Walter Murray, Michael A. Saunders*, and Margaret H. Wright, "Sparse Matrix Methods in Optimization."

10:15 Coffee break.

Wednesday, 11:00 a.m.

Session 12A. Contributed Papers. 3rd Conference Room

Chairperson: Albert M. Guttman

10:40 a.m. Bahram Nour-Omid* and Michael T. Manteig, "Generalized Discretization for Solution of Finite Element Equations.

11:00 A. Jennings* and M. A. Atte, "Incomplete Methods for Solving $A\bar{A}x = b$.

11:20 Tony F. Chan* and Kenneth I. Jackson, "Nonlinearly Preconditioned Krylov Subspace Methods for Discrete-Newton Algorithms.

11:40 Bahram Nour-Omid* and Rencsőerdő G. Parlett, "Element Pre-conditioning."

Session 12B. Contributed Papers. Draft-Bonnerster Room

Chairperson: Margaret H. Wright

10:40 a.m. R. J. Hanson and K. L. Hiebert*, "A Sparse Linear Programming Subprogram."

11:00 Walter Murray, "Null-Space Methods for Large-Scale Quadratic Programming."

11:20 S. Thomas McCormick, "A Fast Algorithm That Makes Matrices Optimally Sparse."

11:40 Michael Engquist, "A Partitioning Approach to Processing Networks and the Solution of Sparse Systems."

12:00 Break.

Session 13A. Contributed Papers. 4th Conference Room

Chairperson: Albert M. Guttman

12:10 p.m. Larry F. Bennett, "Parallel and Adaptive Structured Methods and Their Use in the Solution of Matrix Equations and Least Squares Problems.

12:30 Daniel B. Szyld* and Robert Byersky, "Some Operations on Sparse Matrices."

12:50 Jian-Xin Ren, "Some Generalizations of the Sherman-Morrison and Generalized Sherman-Morrison Formulas for Matrices."

Wednesday, October 27

Session 13B. Contributed Papers. Druid-Dorchester Room

Chairperson: Margaret H. Wright

12:10 p.m. Thomas F. Coleman, "Software for Sparse Matrix Estimation."

12:30 Paul H. Calamai* and Andrew R. Conn, "Continuous ℓ_p Norm Location Problems: A Second-Order Approach."

12:50 Noel E. Courtier, "A New Method for Solving Systems of Linear Inequalities."

1:10 End of Symposium.

Poster Session. Wilshire-Canterbury-Windsor Room

(Posters on display Monday and Tuesday)

D. A. Calahan, "Performance of Sparse Equation Solvers on the CRAY-1."

Iain S. Duff, Roger G. Grimes, John G. Lewis, and W. G. Poole, Jr., "Sparse Matrix Test Problems."

R. E. Funderlic and R. J. Plemmons, "An Incomplete Factorization Method for Singular Irreducible M-Matrices."

James E. Giles, "Implementation of a Large-Scale Linear Programming Package on a Minicomputer."

Roger G. Grimes, John G. Lewis, and William G. Poole, Jr., "Program for the Comparison of Reordering Algorithms for the Solution of Unsymmetric Sparse Systems of Equations."

Fred G. Gustavson, "An Algorithm for Computing a Full Assignment for A Sparse Matrix."

William A. Loden, "A Comparison of Two Sparse Matrix Processing Techniques for Structural Analysis Applications."

TIMETABLES

TIMETABLE

Monday, October 25

Time	St. George Room	Druid-Dorchester Room	Wiltshire-Canterbury-Windsor Room
9:00 - 9:45 a.m.	I. S. Buff*, J. K. Reid		
9:45 - 10:30 a.m.	S. C. Eisenstat		Poster Session
11:00 - 11:20 a.m.	A. M. Grisman, R. G. Grimes, J. A. Lewis*, W. G. Poole	R. L. Cox	D. A. Callahan
11:45 - 12:45 p.m.	A. Drud	G. Sewell	I. S. Buff, R. G. Grimes, J. A. Lewis*, W. G. Poole
1:00 - 1:45 p.m.	R. Fourier	J. W. Neuberger, R. J. Renka*	R. L. Gardner, R. J. Renka*
2:00 - 2:45 p.m.	R. H. Hough, H. F. McAllister, W. J. Stewart*	K. Kaneko	R. L. Gardner, R. J. Renka*
3:00 - 3:45 p.m.	R. V. Graffing, B. R. Snier, D. U. Werner*	M. T. Heath, P. J. Plemmons*, P. G. Ward	R. L. Gardner, R. J. Renka*
4:00 - 4:45 p.m.	L. Kaufman	A. S. Rao	R. L. Gardner, R. J. Renka*
5:00 - 5:45 p.m.	T. Manteuffel*, V. Faber	D. A. Callahan	R. L. Gardner, R. J. Renka*
6:00 - 6:45 p.m.	A. Greenbaum	R. D. Coleman*, E. J. Kushner	A. Sameh*, C. Raft
7:00 - 7:45 p.m.	A. Brandt, S. McCormick*, J. Ruge		R. A. Kington
8:00 - 8:45 p.m.			S. N. Parlett

TIME TABLE

Tuesday, October 26

Time	St. George Room	Druid-Dorchester Room	Wiltshire-Canterbury-Windsor Room
8:45 - 9:30 a.m.	M. T. Heath		Poster Session
9:30 - 10:15 a.m.	A. J. Pope		
		D. A. Callahan	
10:45 - 11:00 a.m.	A. K. Kevorkian*, F. G. Gustavson, G. D. Hachtel	D. S. Scott	
11:00 - 11:20 a.m.	F. G. Gustavson*, G. D. Hachtel, A. K. Kevorkian	J. K. Culicum*, R. A. Willoughby	I. S. Duff, R. G. Grimes, J. G. Lewis, W. G. Poole
11:20 - 11:40 a.m.	G. D. Hachtel*, F. G. Gustavson, A. K. Kevorkian	P. S. Jensen	R. E. Funderlic, R. J. Plemons
11:45 - 12:00 noon	F. G. Gustavson	H. D. Simon	J. E. Giles
12:10 - 12:30 p.m.	S. Ursic	A. George, J. Liu*	
12:30 - 12:50 p.m.	J. de Pillis*, W. Niethammer	A. George, E. Ng*	R. G. Grimes, J. G. Lewis, W. G. Poole
12:50 - 1:10 p.m.	C. R. Johnson	M. S. Kamel*, K. Singhai	F. G. Gustavson
1:30 - 1:50 p.m.	I. S. Duff, J. K. Reid*	A. Behie, P. A. Forsyth*	
1:50 - 3:10 p.m.	J. R. Gilbert*, R. Schreiber	W. I. Bertiger*, D. A. Callahan P. T. Woo	W. A. Loden
3:15 - 3:30 p.m.	A. M. Erisman		W. P. Kamp
3:30 - 4:30 p.m.	W. M. Coughran, W. Fitchner, E. H. Grosse, D. J. Rose*		

TIME TABLE

Wednesday, October 27

Time	St. George Room	Druid-Dorchester Room
8:45 - 9:30 a.m.	D. P. O'Leary	
9:30 - 10:15 a.m.	P. E. Gill, W. Murray, M. A. Saunders*, M. H. Wright	
10:45 - 11:30 a.m.	B. Nour-Omid*, H. D. Simon	R. J. Hanson, R. J. Hether*
11:30 - 12:15 p.m.	A. Jennings*, M. A. Ajiaz	W. Murray
12:45 - 1:30 p.m.	J. F. Chan*, K. Jackson	S. J. McCorrist
1:30 - 2:15 p.m.	B. Nour-Omid*, B. N. Parlett	M. Jagquist
2:15 - 3:00 p.m.	L. F. Bennett	L. F. Thompson
3:00 - 3:45 p.m.	B. E. Szyld*, O. Vishnepolsky	P. A. Calhoun*, A. M. Jones
3:45 - 4:30 p.m.	J.-X. Chen	N. J. Corley

AUTHOR INDEX

AUTHOR INDEX

NAME	ABSTRACT PAGE	SESSION	TIME
Ajiz, M. A.	71	12A	11:00-11:20 a.m. Wednesday
Behie, A.	35	9B	11:30-11:50 a.m. Wednesday
Bennett, Larry F.	36	13A	12:10-12:30 p.m. Wednesday
Bertiger, W. I.	37	9B	11:30-11:50 a.m. Wednesday
Brandt, A.	38	4A	8:00-8:30 a.m. Wednesday
Calahan, D. A.	39	4B	11:30-11:50 a.m. Wednesday
Calahan, D. A.	37	9B	11:30-11:50 a.m. Wednesday
Calahan, D. A.	40	Poster	Monday - Tuesday
Calamai, Paul H.	41	13B	12:30-1:30 p.m. Wednesday
Chan, Tony F.	42	12A	11:30-12:30 p.m. Wednesday
Coleman, Ronald D.	43	4B	11:30-11:50 a.m. Wednesday
Coleman, Thomas F.	44	13B	12:30-1:30 p.m. Wednesday
Conn, Andrew R.	41	13B	12:30-1:30 p.m. Wednesday
Cortey, Noel E.	45	13B	12:30-1:30 p.m. Wednesday
Coughran, William M.	25	10	8:00-8:30 a.m. Wednesday
Cox, R. L.	46	2B	11:30-11:50 a.m. Wednesday
Cullum, Jane	47	14	11:30-11:50 a.m. Wednesday
Dearing, P. M.	48	2B	11:30-11:50 a.m. Wednesday
Den, Jian-Xin	49	13B	12:30-1:30 p.m. Wednesday
de Pillis, John	50	13B	12:30-1:30 p.m. Wednesday
Drud, Arne	51	13B	12:30-1:30 p.m. Wednesday
Duff, Iain S.	26	10	8:00-8:30 a.m. Wednesday
Duff, Iain S.	53	13B	12:30-1:30 p.m. Wednesday
Duff, Iain S.	52	13B	12:30-1:30 p.m. Wednesday
Eisenstat, Stanley C.	27	13B	12:30-1:30 p.m. Wednesday
Engquist, Michael	54	13B	12:30-1:30 p.m. Wednesday
Erisman, Albert M.	55	13B	12:30-1:30 p.m. Wednesday
Erisman, Albert M.	56	13B	12:30-1:30 p.m. Wednesday
Faber, Vance	57	13B	12:30-1:30 p.m. Wednesday
Fichtner, Wolfgang	25	13B	12:30-1:30 p.m. Wednesday
Forsyth, P. A.	35	13B	12:30-1:30 p.m. Wednesday
Fourer, Robert	57	13B	12:30-1:30 p.m. Wednesday
Funderlic, R. F.	58	13B	12:30-1:30 p.m. Wednesday

NAME	ABSTRACT PAGE	SESSION	TIME
George, Alan	59	8B	12:10-12:30 p.m. / Tues.
George, Alan	60	8B	12:30-12:50 p.m. / Tues.
Gilbert, John R.	51	9A	7:50- 8:10 p.m. / Tues.
Giles, James L.	62	Poster	Monday - Tuesday
Gill, Philip E.	28	11	9:30-10:15 a.m. / Wed.
Greenbaum, Anne	63	4A	7:50- 8:10 p.m. / Mon.
James, Roger G.	56	2A	11:00-11:20 a.m. / Mon.
James, Roger G.	52	Poster	Monday - Tuesday
James, Roger G.	64	Poster	Monday - Tuesday
Jousse, Eric H.	25	10	8:45- 9:30 p.m. / Tues.
Justavson, Fred G.	78	7A	10:40-11:00 a.m. / Tues.
Justavson, Fred G.	67	7A	11:00-11:20 a.m. / Tues.
Justavson, Fred G.	68	7A	11:20-11:40 a.m. / Tues.
Justavson, Fred G.	69	7A	11:40-12:00 noon / Tues.
Justavson, Fred G.	66	Poster	Monday - Tuesday
Kaftan, Barry J.	74	7A	10:40-11:00 a.m. / Tues.
Kaftan, Barry J.	67	7A	11:00-11:20 a.m. / Tues.
Kaftan, Barry J.	68	7A	11:20-11:40 a.m. / Tues.
Kansin, R. J.	53	12B	10:40-11:00 a.m. / Wed.
Heath, Michael T.	79	3B	12:30-12:50 p.m. / Mon.
Heath, Michael T.	77	6	8:45- 9:30 a.m. / Tues.
Heppert, K. L.	75	2B	10:40-11:00 a.m. / wed.
Jackson, Kenneth	70	12B	11:20-11:40 a.m. / wed.
Jennings, A.	71	2A	11:00-11:20 a.m. / wed.
Jensen, Paul S.	72	1B	11:20-11:40 a.m. / Tues.
Johnson, C. R.	73	1A	12:30- 1:15 p.m. / Tues.
Kamel, M. S.	74	10	10:40-11:00 a.m. / Tues.
Kamp, W. P.	75	10	10:40-11:00 a.m. / Tues.
Kaneko, Ken	76	10	10:40-11:00 a.m. / Mon.
Kaufman, Linda	77	10	10:40-11:00 a.m. / Mon.
Kevorkian, Aram A.	78	10	10:40-11:00 a.m. / Tues.
Kevorkian, Aram A.	79	10	10:40-11:00 a.m. / Tues.
Kevorkian, Aram A.	80	10	10:40-11:00 a.m. / Tues.
Kilby, Robert	71	10	10:40-11:00 a.m. / Mon.
Kishimoto, Tetsuro	72	10	10:40-11:00 a.m. / Mon.

NAME	ABSTRACT PAGE	SESSION	TIME
Lewis, John G.	56	2A	11:00-11:20 a.m. / Mon.
Lewis, John G.	52	Poster	Monday - Tuesday
Lewis, John G.	64	Poster	Monday - Tuesday
Liu, Joseph	59	8B	12:10-12:30 p.m. / Tues.
Loden, William A.	80	Poster	Monday - Tuesday
Manteuffel, Tom	81	4A	7:30- 7:50 p.m. / Mon.
McAllister, David F.	79	3A	12:10-12:30 p.m. / Mon.
McCormick, S.	38	4A	8:10- 8:30 p.m. / Mon.
McCormick, S. Thomas	82	12B	11:20-11:40 a.m. / Wed.
Murray, Walter	28	11	9:30-10:15 a.m. / Wed.
Murray, Walter	83	12B	11:00-11:20 a.m. / Wed.
Neuberger, J. W.	84	2B	11:40-12:00 noon / Mon.
Ng, Esmond	60	8B	12:30-12:50 p.m. / Tues.
Niethammer, Wilhelm	50	8A	12:30-12:50 p.m. / Tues.
Nour-Omid, Bahram	86	12A	10:40-11:00 a.m. / Wed.
Nour-Omid, Bahram	85	12A	11:40-12:00 noon / Wed.
O'Leary, Dianne P.	30	11	8:45- 9:30 a.m. / Wed.
Parlett, Beresford N.	31	5	8:45- 9:30 p.m. / Mon.
Parlett, Beresford N.	85	12A	11:40-12:00 noon / Wed.
Plemonns, R. J.	70	3B	12:30-12:50 p.m. / Mon.
Plemonns, R. J.	58	Poster	Monday - Tuesday
Poole, William G.	52	Poster	Monday - Tuesday
Poole, William G.	56	2A	11:00-11:20 a.m. / Mon.
Poole, William G.	64	Poster	Monday - Tuesday
Pope, Allen J.	32	6	9:30-10:15 a.m. / Tues.
Rao, A. S.	87	3B	12:50- 1:10 p.m. / Mon.
Reid, John K.	26	1	9:00- 9:45 a.m. / Mon.
Reid, John K.	53	9A	7:30- 7:50 p.m. / Tues.
Renka, R. J.	84	2B	11:40-12:00 noon / Mon.
Rose, Donald J.	25	10	8:45- 9:30 p.m. / Tues.
Ruge, J.	38	4A	8:10- 8:30 p.m. / Mon.

NAME	ABSTRACT PAGE	SESSION	TIME
Sameh, A.	88	4B	8:10- 8:30 p.m. / Mon.
Saunders, Michael A.	28	11	9:30-10:15 a.m. / Wed.
Schreiber, Robert	61	9A	7:50- 8:10 p.m. / Tues.
Scott, David S.	89	7B	10:40-11:00 a.m. / Tues.
Sewell, Granville	90	2B	11:20-11:40 a.m. / Mon.
Shier, D. R.	48	3A	12:30-12:50 p.m. / Mon.
Simon, Horst D.	91	7B	11:40-12:00 noon / Tues.
Simon, Horst D.	86	12A	10:40-11:00 a.m. / Wed.
Singhal, K.	74	8B	12:50- 1:10 p.m. / Tues.
Stewart, William J.	79	3A	12:10-12:30 p.m. / Mon.
Szyld, Daniel B.	92	13A	12:30-12:50 p.m. / Wed.
Taft, C.	88	4B	8:10- 8:30 p.m. / Mon.
Ursic, Silvio	93	8A	12:10-12:30 p.m. / Tues.
Vishnepolsky, Oleg	92	13A	12:30-12:50 p.m. / Wed.
Ward, R. C.	70	3B	12:30-12:50 p.m. / Mon.
Warner, D. D.	48	3A	12:30-12:50 p.m. / Mon.
Willoughby, Ralph A.	47	7B	11:00-11:20 a.m. / Tues.
Wright, Margaret H.	28	11	9:30-10:15 a.m. / Wed.
Woo, P. T.	37	9B	7:50- 8:10 p.m. / Tues.

INVITED ABSTRACTS

NUMERICAL SIMULATION OF VLSI CIRCUITS

William M. Coughran, Jr.
Wolfgang Fichtner
Eric H. Grosse
Donald J. Rose*
Bell Laboratories
Murray Hill, New Jersey 07974

Circuit analysis has motivated interesting numerical analysis for several decades. We will review previous work with particular emphasis on the type and structure of the nonlinear and linear equations that arise. We then suggest a new formulation of the circuit equations based on a hierarchical structuring of the circuit. Our reformulation of the equations leads naturally to the circuit substructuring that is often termed "macromodeling", currently in vogue.

We will present several numerical examples of macromodeling. In particular we will emphasize the construction of a transistor device macromodel directly from a device simulator as well as from the more traditional equivalent circuit.

Session 10. 8:45 - 9.30 p.m., Tuesday

*Indicates the presenter of the paper, if more than one author is given.

DIRECT METHODS FOR SOLVING
SPARSE SYSTEMS OF LINEAR EQUATIONS

Iain S. Duff*
John K. Reid
AERE Harwell
Oxon, OX11 ORA, ENGLAND

We survey algorithms and software for solving sparse systems of linear equations paying particular attention to recent developments. We classify the various algorithms according to the type of system they solve (i.e., unsymmetric, symmetric definite, symmetric indefinite, unsymmetric but with symmetric pattern) and whether they perform pivoting for numerical stability. We consider both algorithms which work in main memory and those which use auxiliary storage.

We illustrate the performance of most of the software we discuss by runs on test problems and give a critical evaluation of each, stressing their strengths, weaknesses and restrictions.

Session 1. 9:00 - 9:45 a.m., Monday

ITERATIVE METHODS FOR SOLVING LARGE SPARSE LINEAR SYSTEMS

Stanley C. Eisenstat
Yale University
New Haven, Connecticut 06520

This talk will (attempt to) survey the state of the art in iterative and semi-iterative methods for solving large sparse systems of linear equations $Ax = b$. Three classes of problems will be considered, corresponding to the coefficient matrix being symmetric positive definite, symmetric indefinite, or nonsymmetric. For each class, the emphasis will be on:

- Krylov subspace methods - which choose the m^{th} iterate from the Krylov subspace $\text{Span}\{b, Ab, A^2b, \dots, A^{m-1}b\}$ and thus only require the ability to multiply a vector by the coefficient matrix (e.g., the conjugate gradient and Chebyshev semi-iterative methods);
- preconditioning techniques - methods for "rescaling" the original system to increase the rate of convergence (e.g., incomplete factorization);
- reduction techniques - methods for reducing a problem in one class to a smaller problem in the same class (e.g., cyclic reduction) or a problem in another class (e.g., forming the normal equations) while maintaining whatever sparseness, symmetry, and definiteness was present in the original system;
- general-purpose software.

The written version of the talk will include an extensive bibliography of recent papers.

Session 1. 9:45 - 10:30 a.m., Monday

SPARSE MATRIX METHODS IN OPTIMIZATION

Philip E. Gill
Walter Murray
Michael A. Saunders*
Margaret H. Wright
Stanford University
Stanford, California 94305

Optimization algorithms typically require the solution of many systems of linear equations $B_k y_k = b_k$. When large numbers of variables or constraints are present, these linear systems could account for much of the total computation time.

Both direct and iterative equation solvers are needed in practice. Unfortunately, most off-the-shelf solvers are designed for single systems, whereas optimization problems give rise to hundreds or thousands of systems. To avoid refactorizing, or to speed the convergence of an iterative method, it is essential to note that B_k is related to B_{k-1} .

We review various sparse matrices that arise in optimization, and discuss compromises that are currently being made in dealing with them. Since significant advances continue to be made with single-system solvers, we give special attention to methods that allow such solvers to be used directly on modified systems (e.g., the PFI update; use of the Schur complement). At the same time, we hope that future improvements to linear-equation software will be oriented more specifically to the case of related matrices B_k .

Session 11. 9:30 - 10:15 a.m., Wednesday

NUMERICAL METHODS FOR LARGE SPARSE LINEAR LEAST SQUARES PROBLEMS

Michael T. Heath
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

Large sparse least squares problems arise in many applications, including geodetic network adjustments and finite element structural analysis. Although geodesists and engineers have been solving such problems for years, it is only relatively recently that numerical analysts have turned attention to them. In this talk we present a survey of numerical methods for large sparse linear least squares problems, focusing mainly on developments since the last comprehensive surveys of the subject published in 1976. We consider direct methods based on elimination and on orthogonalization, as well as various iterative methods. The ramifications of rank deficiency, constraints, and updating are also discussed.

Session 6. 8:45 - 9:30 a.m., Tuesday

SOLVING MESH PROBLEMS ON PARALLEL PROCESSORS

Marine R. McLeary
University of Maryland
College Park, Maryland 20742

Suppose that the sparsity structure of an optimization problem or system of linear or nonlinear equations corresponds to a rectangular grid of nodes in which each node is linked to any or all of its north, south, east, west, northeast, northwest, southeast, and southwest neighbors. Such problems arise in fields including discretization of partial differential equations, network problems, and image processing. If the problem is to be solved on a parallel computer using an iterative method, the Jacobi method allows parallelism, but special orderings of nodes are required to exploit parallelism in methods such as Gauss-Seidel, SOR, or conjugate gradients.

In this work, we discuss requirements on parallel computer architectures and algorithms which permit efficient solution of large sparse problems. Node orderings and processor arrangements are presented which assign each processor to a small number of points and enable each to work in parallel with other processors and only limited data transfer among processors. Examples discussed include the nine-point finite difference operator and an optimization problem with bound constraints.

Session 11. 8:45 - 9:30 a.m., Wednesday

SOFTWARE FOR SPARSE EIGENVALUE PROBLEMS

Berestford N. Parlett
University of California
Berkeley, California 94720

There are several good packages for solving sparse linear systems, yet nothing comparable for eigenvalue calculations. Why not?

This question will be discussed along with a survey of what is available to the public. Some good software is buried in applications packages in engineering, chemistry, and physics centers.

Some judgments and comparisons will be offered concerning subspace iteration and versions of the Lanczos algorithm. Of more importance is improved understanding, and consensus, among the experts since 1978.

It seems as though eigenvalue programs will have to be specially crafted for computers such as the Cray 1. This prospect poses the question of the right level of portability to contemplate for software which is going to make more demands on the host system than does any subroutine in EISPACK. The use of secondary storage is a key factor in the efficiency of a program and this makes it difficult to write software which is both effective and independent of the operating system.

Session 5. 8:45 - 9:30 p.m., Monday

GEODETIC COMPUTATIONS AND SPARSITY

Allen J. Pope
National Geodetic Survey
NOAA/NOS
Rockville, Maryland 20852

The least squares adjustment of geodetic networks is a natural arena for the application of sparse matrix technology. In this talk, geodetic computations are reviewed for the non-geodesist. Special emphasis is given to those features (of which there are many) that set these computations apart from most other sparse problems. A survey of the geodetic involvement with sparsity begins near the origins of the science and continues through recent trials of various re-ordering strategies. Although it can be maintained that the exploitation of sparsity has reached a certain plateau, there still remain numerous challenges. In particular these arise from the large size of some geodetic data sets, such as that involved in the re-adjustment of the North American datum.

Session 6. 9:30 - 10:15 a.m., Tuesday

CONTRIBUTED ABSTRACTS

INCOMPLETE FACTORIZATION METHODS FOR FULLY IMPLICIT SIMULATION OF ENHANCED OIL RECOVERY

A. Benie
P. A. Forsyth Jr.*
Computer Modelling Group
Calgary, Alberta, Canada T2L 2A6

Fully implicit simulation of enhanced oil recovery using thermal (steam and in situ combustion) methods gives rise to a highly structured block-banded non-symmetric Jacobian. Solution of large problems requires effective iterative methods.

Recently, incomplete factorization methods (ILU) have been used to solve these systems. In this paper, several ILU methods are developed using natural, diagonal (D2) and alternating diagonal (D4) ordering. Diagonal ordering was first used for symmetric systems by Watts, while Tan and Letkeman suggested alternate diagonal ordering for non-symmetric problems. The method of Watts is generalized to the strongly non-symmetric case, and an improvement to the algorithm of Tan and Letkeman is developed which saves a considerable amount of work. In addition, several degrees of factorization are used for all these orderings. These techniques are all accelerated with the ORTHOMIN algorithm.

Each of these methods is developed with vector machines in mind, and particular attention is paid to those portions of the algorithm which can be readily vectorized. All the ILU methods can also be used with the COMBINATIVE technique.

Results are presented for several model problems and test Jacobians generated from steam simulations. The results show the effects of:

- (1) Natural, D2 and D4 ordering in two and three dimensions.
- (2) Varying the degree of the factorization.
- (3) Use of the modified factorization.
- (4) The effectiveness of these methods in scalar and vector mode on the CRAY-1.

Session 9B. 7:30 - 7:50 p.m., Tuesday

*Indicates the presenter of the paper, if not the author.

ACCELERATED ITERATIVE PROJECTION METHODS AND THEIR USE IN THE SOLUTION OF MATRIX EQUATIONS AND LEAST SQUARES PROBLEMS

John R. Bunch
University of Colorado
Boulder, Colorado 80302

A very general theory of accelerated iterative projection methods with relaxation factors is presented. The theory includes accelerated refined, fine, coarse, fine, and vectors on finite dimensional subspaces of \mathbb{C}^m or \mathbb{R}^m , accelerated generalized column projection methods (called Right Projective Iterative Methods) for the solution of $m \times n$ matrix equations and linear least squares problems, and accelerated generalized row projection methods (left Projective Iterative Methods) for the solution of $m \times n$ matrix equations and linear least squares problems. All projective methods discussed are finite methods, so that like the method of conjugate gradients, they are guaranteed to converge to the exact solution sought in a finite number of steps if infinite precision arithmetic were possible. Algorithms and software for computerized implementation of the methods in solving sparse matrix equations and least squares problems as well as operation counts are mentioned. Important properties of the new methods, such as those which follow are stressed. The methods appear to work comparatively well on any matrix least squares problem or matrix equation of the form $Ax = b$, where A is any complex or real-valued $m \times n$ matrix and b is a complex or real-valued m -tuple. No special hypothesis on the structure of the matrix A is needed. For large scale matrix problems, convergence to an acceptable solution often occurs much sooner than predicted for the exact solution. The methods can produce extremely accurate results, and in the case of least squares problems, do not appear to increase ill-conditioning as normal equations sometimes do. The methods may be further accelerated in some cases by using relaxation factors. In many cases, single precision arithmetic may be employed to generate accurate solutions, even in the case of ill-conditioned systems. The rate of convergence for the new methods often increases as iterations progress. Storage space required is relatively small, with matrices stored using row-wise packing or column-wise packing, depending on the method of solution to be employed. If preferred, rows or columns of the matrix may be generated as needed.

In conclusion, results obtained by applying the methods to test problems are presented. Applications to 34×34 and 167×167 unsymmetric ill-conditioned matrix equations are mentioned. Results obtained in applying the methods to solve symmetric 31×31 and 361×361 matrix equations generated in attempting to approximate the solution of a boundary value problem using finite difference equations are given. A number of attempts to solve a couple of problems included in Iain S. Duff's collection of sparse matrix problems are discussed. Specifically, an unsymmetric matrix equation generated by a matrix obtained from a laser printer and a 10×10 linear least squares problem is considered.

Session 13A, 12:10 - 12:30 p.m., Wednesday

THE EFFECT OF COMPUTER ARCHITECTURE ON DIRECT SPARSE MATRIX ROUTINES IN PETROLEUM RESERVOIR SIMULATION

W. I. Bertiger*

Chevron Oil Field Research Company
La Habra, California 90631

D. A. Calahan

University of Michigan
Ann Arbor, Michigan 48109

P. T. Woo

Chevron Oil Field Research Company
La Habra, California 90631

Large systems of linear equations are solved daily in reservoir simulation. These equations arise from finite difference approximations to systems of partial differential equations describing the flow of fluids in a reservoir. With the advent of new vector computer technology the algorithms for solving the linear equations must be changed to make better use of the machine architecture.

In solving linear equations on a vector computer, there is a trade off between computation rate and work (the amount of arithmetic). At this stage of development, optimum use of the vector computer has led us to change from a Yale general purpose sparse matrix routine with the inner loops coded in assembly language to an assembly coded bandsolver developed at the University of Michigan. The sparse routine is optimized to do the least amount of arithmetic and is best suited for a scalar computer. The bandsolver is, on the other hand, optimized to achieve maximum computation rate on a vector computer. The trade off for solving linear equations in reservoir simulation is in favor of the bandsolver.

When equations involving the wellbore pressures are added to the flow equations, the added portion of the matrix is not banded. We will describe how the bandsolver can still be used effectively with an algorithm attributed to George. Comparisons between the sparse and bandsolvers based on operation counts and actual timing runs on the Cray-1b will be presented for model and real reservoir problems. Results for the Cray XMP will be presented if the benchmark runs can be completed in time. Future improvements to the direct solution of linear equations on vector computers will be discussed.

Session 9B. 7:50 - 8:10 p.m., Tuesday

ALGEBRAIC MULTIGRID (AMG) FOR GEODETIC EQUATIONS

A. Brandt
Weizmann Institute of Science
Rehovot, Israel

S. McCormick*
J. Ruge
Colorado State University
Fort Collins, Colorado 80523

Multigrid algorithm design requires that its major processes be tailored to each application. This can often be done manually at the time the discretization procedure is being designed, but not always. The engineer may, for example, preselect a fine grid finite element triangulation so that determining even what the coarser grids should be is difficult if not impossible. Moreover, some problems, such as those that arise in geodetic surveying, are inherently discrete. So it is important to consider methods for automatic multigrid design.

Algebraic multigrid (AMG) is one such method, where the discrete problem, represented either by a matrix or operator stencil, is subjected to a preprocessing stage. AMG bases its decisions on the concept of strong coupling which, loosely speaking, is a way to interpret effective dependencies between variables via coefficients in the discrete operator that connect them. This concept provides a motivation for deciding which variables are to constitute the coarse "grid" and for determining interpolation. The remaining multigrid processes can be developed from this design in a fairly straightforward way.

This is a preliminary report on experience with several versions of AMG applied to elliptic differential problems that include diffusion equations (with coefficients that are both widely varying and strongly discontinuous), anisotropic fluid flows, and non-uniform grids. Special attention is given to the potential of AMG for solving purely algebraic problems such as the large-scale geodetic equations.

Session 4A, 8:10 - 8:30 p.m., Monday

SPARSE DIRECT METHODS FOR VECTOR MULTIPROCESSORS

D. A. Calahan
University of Michigan
Ann Arbor, Michigan 48109

The CRAY X-MP is a harbinger of a class of forthcoming GIGAFLOP supercomputers composed of arrays of vector processors connected to a common main memory. Related sparse matrix methods must combine the parallel (multiprocessor) and vector taxonomies that have been studied independently in the past.

Computationally, the solution may be phrased as a two-level synchronization.

1. The vectorized inner loop, precisely synchronized at the instruction level, and operating independently among the processors.
2. The parallel outer loop, approximately synchronized at the instruction block level.

If the vectors are sufficiently long, both loops may be precisely synchronized by distributing the inner loop vectors among the processors (outer loop) and then having them operate in an SIMD mode. This is undoubtedly the manner in which large full and banded systems will be solved, but is algorithmically uninteresting.

Rather, consider the case of a randomly-sparse matrix where each element is a rectangular block representing coupling of unknowns at two grid points in a finite element or finite difference representation. Further, these blocks are different sizes corresponding to the different number of unknowns at each grid point, and may change dynamically in a time-dependent problem.

It is proposed that a solution proceed by

1. Local decoupling of blocks by reordering of elements (consistent, for example, with nested dissection), and
2. Dynamic scheduling of independent block processing among the processors; dynamic scheduling is necessary to compensate for different block sizes.

An extended CRAY X-MP architecture is examined to project performance of such an algorithm. A simulator is being developed for this machine, and will quite likely permit some precise projections by the time of the conference. Actual runs are also possible.

Session 4B. 7:30 - 7:50 p.m., Monday

PERFORMANCE OF SPARSE EQUATION SOLVERS ON THE CRAY-1

D. A. Cullum
University of Michigan
Ann Arbor, Michigan 48109

Algorithms and performance of five classes of sparse matrix solvers are discussed.

1. General sparse solvers. These solve randomly-sparse equations described symbolically in column-ordered form and are compatible with traditional scalar packages. The approaches investigated both involve two-step symbolic/numeric processing (in the manner of Gustavson, et al).

- (a) Hybrid vector-scalar solution of medium-density systems. The symbolic phase identifies dense segments of columns with more than a user-prescribed value (n) of contiguous non-zeros. These segments are processed in vector mode and the remainder in scalar mode. Maximum rate: 37 MFLOPS.
- (b) Decoupled scalar solution of highly-sparse systems. A variation of code generation methods for a pipelined processor involves the local decoupling of highly-sparse equations (in the manner of nested dissection). Scalar machine code is then generated which seeks to "cram" the floating point pipelines with independent scalar code. Typical rate: 15 MFLOPS.

2. Special sparse solvers. Higher performance can be achieved when special sparsity structure can be identified by the user. In general, vectors can be usefully defined either within a dense submatrix or across similarly-structured submatrices.

- (a) General block-oriented solver. From a block-oriented input matrix description, dense rectangular submatrices may be processed in blocks to reduce indirect addressing overhead and to exploit the CRAY-1 register/main-memory hierarchy. Maximum rate: 15 MFLOPS.
- (b) Banded and profile solvers. Efficiently-coded symmetric and unsymmetric banded and profile solvers achieve high execution rates with relatively small bandwidths, but in order to profit from blocking, bandwidths must be limited to 100-300 columns (half bandwidths, in turn, from 50-150).
- (c) Structured matrix solvers. When solving coupled subsystems with identical structure, one can exploit the parallel electronic architecture. One can solve the subsystems simultaneously, in fact, using 100-150 MFLOPS.

CONTINUOUS ℓ_p NORM LOCATION PROBLEMS: A SECOND-ORDER APPROACH

Paul H. Calamai*
Andrew R. Conn
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

A second-order algorithm for solving a prototype continuous min-sum multifacility location problem involving ℓ_p distances is presented. This problem corresponds to a special case of the p -median problem and a continuous version of the quadratic assignment problem.

We demonstrate how projectors can be used to circumvent the difficulties associated with the nondifferentiability of these problems. The technique, which is an extension of an earlier first-order method used by the author, provides a unified, numerically sound and stable approach. The implementation is the first to efficiently exploit the special structure of the problems under consideration. In particular, 1) we are able to solve the linear systems that arise in the development by stable means that do not suffer from fill-in, 2) a special linesearch which is based, in part, on the results presented in a paper by Overton recognizes the possibility of first derivative discontinuities and second derivative unboundedness along descent directions, and 3) the degeneracies which are an inherent characteristic of these problems are dealt with using simple perturbations.

Although details are initially given for only unconstrained fixed ℓ_p norm problems we show how the framework can be readily extended to mixed norm problems as well as to constrained problems.

Session 13B. 12:30 - 12:50 p.m., Wednesday

NONLINEARLY PRECONDITIONED KRYLOV SUBSPACE METHODS FOR DISCRETE-NEWTON ALGORITHMS

Tony F. Chan*
Yale University
New Haven, Connecticut 06511

Barry F. Saad
University of Toronto
Toronto, Ontario, Canada M5S 1A1

One promising idea that some people have had in the application of Krylov subspace methods for solving the linear systems $\mathbf{Aw} = \mathbf{b}$ that arise in the inner loop of a Newton-like iterative method for solving nonlinear systems $\mathbf{F}(\mathbf{x}) = \mathbf{0}$ (with \mathbf{A} being the Jacobian of \mathbf{F}) is the use of the directional differencing $\mathbf{F}(\mathbf{x} + \mathbf{d}) - \mathbf{F}(\mathbf{x})$ for approximating the matrix-vector product \mathbf{Au} in the Krylov subspace methods. This requires only function evaluations and avoids explicit evaluation and storage of the Jacobian. We are interested in accelerating the convergence rate of the inner loop by preconditioning \mathbf{A} , while still employing directional differencing. However, since most preconditioning techniques are derived from the matrix elements of \mathbf{A} explicitly, it is not obvious how to apply them when \mathbf{A} is not explicitly available. We have derived an algorithm for preconditioning the Krylov subspace method with directional differencing that reduces to the SSOR preconditioning of \mathbf{A} in the linear case. It only requires evaluating the diagonal elements of the Jacobian, which can be approximated by function evaluations. The overall algorithm is thus well-suited for large and sparse problems, especially when function evaluations are not too expensive. Numerical experiments show that the nonlinear preconditioning is as effective as in the linear case.

Session 12A. 11:20 - 11:40 a.m., Wednesday

THE SPARSE MATRIX LIBRARY FOR THE FPS-164 ATTACHED PROCESSOR

Ronald D. Coleman*
Edward J. Kushner
Floating Point Systems, Inc.
Portland, Oregon 97223

This paper describes the sparse matrix routines for the Floating Point Systems FPS-164 Attached Processor and presents associated benchmark results. The highly parallel and pipelined architecture of the FPS-164 provides high performance for vector and matrix computations at relatively low costs. Included as part of the FPS-164 Program Development Software is the APMATH64 Math Library which contains nearly 500 subroutines that are organized into 14 sub-libraries. One sub-library is the Sparse Matrix Library which contains 13 routines for the solution of sparse linear systems by both direct and iterative methods. The Sparse Matrix Library also includes 12 routines for performing matrix arithmetic operations with sparse matrices. In addition to these routines, the Advanced Math Library (also part of APMATH64) includes a sparse linear programming routine and a profile oriented linear system solver routine. The library routines are written in APAL64, the FPS-164 assembly language. Typically, these routines execute 15 to 25 times faster than the equivalent FORTRAN routines on the VAX-11/780.

Session 4B. 7:50 - 8:10 p.m., Monday

SOFTWARE FOR SPARSE MATRIX ESTIMATION

Thomas F. Coleman
Cornell University
Ithaca, New York 14853

The solution of sparse systems of nonlinear equations or differential equations usually requires the determination of the Jacobian matrix of a nonlinear mapping. Often it is advantageous to estimate the Jacobian matrix by function (or perhaps gradient) differences. When the Jacobian (or perhaps Hessian) matrix has a known sparsity structure, the number of function evaluations needed for the estimation can be quite small compared to the dimension of the problem. Recently, Coleman and Moré have exploited a graph theoretic view to develop efficient algorithms to estimate sparse Jacobian and Hessian matrices. In this talk we describe the resulting software that has been developed, at Argonne National Laboratory, for the sparse matrix estimation problem.

Session 13B. 12:10 - 12:30 p.m., Wednesday

A NEW METHOD FOR SOLVING SYSTEMS OF LINEAR INEQUALITIES

Noel E. Courtney
Lockheed-California Company
Burbank, California 91520

A new iterative descent method to perform least-squares is presented. It is shown how a local weighted least-squares solution can be achieved by adding extra variables and extra equations to a system of polynomial equations and performing sequential linear weighted least-squares. When this new method is applied to solving systems of linear inequalities it always leads to a solution when there is one. When there is no solution it converges to an approximate least squares solution. Since the convergence rate depends on the initial guess it is difficult to compare this method with others. Nevertheless it is established that globally the convergence rate is not an increasing function of $m := \min(e, v)$ where e is the number of equations and inequalities and v is the number of variables. Of course this method is a good candidate to take advantage of existing sparse matrix algorithms for achieving linear least-squares solutions.

Session 13B. 12:50 - 1:10 p.m., Wednesday

**LSOD28: A VARIANT OF LSODE FOR PROBLEMS HAVING A
GENERAL LARGE SPARSE JACOBIAN**

R. L. Cox
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

LSOD28 represents a merging of LSODE (a state-of-the-art computer program written by A. C. Hindmarsh for numerically integrating stiff systems of first order ordinary differential equations) and MA28 (a package written by I. S. Duff for solving sets of sparse unsymmetric linear equations) which is designed to handle stiff problems in which the Jacobian is a general large sparse unsymmetric matrix. Because MA28 performs LU decompositions incorporating partial pivoting for numerical stability, LSOD28 avoids the requirement that the Jacobian be positive definite or diagonally dominant as demanded by some previously existing codes which do not incorporate pivoting.

The use of sparse matrix techniques as opposed to manipulation of the full matrix results in great savings in both computer storage and execution time for large problems. At the same time, questions must be faced which are of little or no concern in the use of full matrix algorithms.

These issues include an efficient user-oriented method for representing the structure (location of nonzeros) of the Jacobian matrix, the frequency at which the pivot strategy is to be revised, the possibility of changes in the structure of the Jacobian (zero derivatives becoming nonzero) during the integration, and the efficient generation of the Jacobian when its elements must be calculated numerically.

Session 2B. 11:00 - 11:20 a.m., Monday

AN ACCELERATED BLOCK LANCZOS PROCEDURE FOR EXTREME EIGENVALUES OF SYMMETRIC MATRICES

Jane Cullum*
Ralph A. Willoughby
IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

Typically in an iterative block Lanczos procedure, the user specifies a priori the amount of computer storage available for the block computation. Therefore, there is a tradeoff between the number of vectors included in the first block and the number of blocks that can be generated on a given iteration. The size of the first block determines the gap between the desired eigenvalues and the eigenvalues not being approximated by the procedure. The more vectors used in the first block the larger these gaps. The number of blocks used within a given iteration determines the effective spread. The more blocks that are used, the smaller the effective spread. The convergence rate depends upon the ratio of these gaps and the effective spread.

Since the user typically doesn't know the gaps a priori, generally speaking the best overall approach is to try to maximize the number of blocks used on each iteration. The number of blocks used in each iteration in the block procedure in Cullum and Donath [1974] could vary as the computations proceeded. The second block was computed using a modified Gram-Schmidt orthogonalization with pivoting for size and any vector whose norm was less than a given program-generated tolerance was dropped from the block. This reduction in size of the second block could lead to an increase in the number of blocks generated on subsequent iterations and an acceleration of convergence. Such an acceleration of convergence however, might not be observed until after a large number of iterations and in fact might never happen due to the particular sizes of the blocks involved.

In this talk we present an iterative block Lanczos procedure that maximizes the number of blocks generated on each iteration. The full unnormalized second block is computed on each iteration. Then, however, using a selection procedure based upon an optimization argument, only the 'best' vector in the second block is allowed to perpetuate itself. Succeeding blocks are reorthogonalized w.r.t. the ancestors of the vectors dropped from the second block. In many cases this approach yields greatly accelerated convergence over what is obtained by the approach in Cullum and Donath [1974]. Large numerical examples are included.

Session 7B. 11:00 - 11:20 a.m., Tuesday

MAXIMAL CHORDAL SUBGRAPHS
AND
DERIVED MATRIX SPLITTINGS

P. M. Dearing
D. R. Shier
D. D. Warner*
Clemson University
Clemson, South Carolina 29631

A linear system of equations has a perfect elimination ordering if and only if the corresponding adjacency graph is chordal. For systems which do not possess a perfect elimination ordering it is customary to try to find an ordering which minimizes the fill-in. One such approach is to embed the corresponding adjacency graph in a minimal chordal supergraph.

Another approach for solving the linear system is to consider an iterative method based on a splitting which is defined by a maximal chordal subgraph. In this paper we present an algorithm for finding a maximal chordal subgraph and explore the effectiveness of the derived iterative scheme.

Session 3A. 12:30 - 12:50 p.m., Monday

SOME EXPERIENCES OF SOLVING LARGE SCALE
GENERALIZED EIGENVALUE PROBLEMS BY LANCZOS METHOD

Qian-Xin Jin
Computer Center of Academia Sinica
People's Republic of China

(Only title was submitted.)

Session 13A. 12:50 - 1:10 p.m., Wednesday

**A STATIONARY ITERATIVE METHOD THAT WORKS EVEN FOR
HERMITIAN $Ax = b$**

John de Pillis*
University of California
Riverside, California 92521

Wilhelm Niethammer
Universitat Karlsruhe
D-7500 Karlsruhe, GERMANY

We develop a theory for finding scalars $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ so as to produce the sequence defined by

$$y_{k+2} = (\alpha_1 B + \alpha_2) y_{k+1} + (\alpha_3 B + \alpha_4) + (\alpha_1 + \alpha_3) A_0^{-1} b, \quad k = 0, 1, 2, \dots$$

where y_0, y_1 are arbitrary. $A = A_0(I - B)$ is an $m \times m$ non-singular matrix and matrix A_0 is easy to invert. The α_i 's are chosen so that $y_k \rightarrow x$, and $Ax = b$. Heretofore, the theory of (stationary) one and two-part sequences, e.g., Jacobi, SOR methods, Chebyschev semi-iterative method (which is asymptotically stationary), required $\sigma(B)$, the spectrum of B , to lie wholly on one side or the other of a line passing through the point $z = 1$. For example, if A is positive definite, $A_0 = I$, $\sigma(B) \subset (-\infty, 1)$. But if A is hermitian and $A_0 = I$, $\sigma(B) \subset (1-t, 1+t)$, and hence straddles the point $z = 1$.

As a special case of our theory, we have the theorem:
Let $A = A^*$ with spectrum $\sigma(A) \subset (-t, t)$. Then for any y_0, y_1 , the sequence

$$y_{k+2} = [A(-y_{k+1} + y_k/2) + b/2]/iR + y_{k+1}$$

always converges to x , where $Ax = b$. The asymptotic rate of convergence is $R_A = -\log(r)$ where $r^2 = (1 + (1-K^{-2})^{1/2})/2$ and K is the condition number of A .

Session 8A. 12:30 - 12:50 p.m., Tuesday

SPIKE SELECTION FOR LARGE SPARSE SETS OF NONLINEAR EQUATIONS

Arne Drud
World Bank
Washington, D.C. 20433

We consider square nonlinear systems with the following properties:

- a. The largest simultaneous block is large (500-3000 equations).
- b. The system is sparse (3-5 nonzero derivatives per equation on average).
- c. Most equations are analytically invertible w.r.t. most of its variables.

Each simultaneous block of such a system is often solved with a spike selection/partitioning/reduction procedure: Select a set of spike variables and spike equations such that the remaining system is recursive and analytically solvable for fixed spike variables. Use a "small scale" algorithm to solve the reduced system, i.e., to solve the spike equations w.r.t. the spike variables, implicitly substituting out all recursive variables.

The criteria for spike selection are to minimize the size of the reduced system and at the same time keep it well-conditioned.

The paper will compare two spike selection algorithms, one based on Hellerman-Rarick with post-processing to eliminate noninvertible relationships in the recursive system (ref. swapping of triangular columns in linear systems), and one that considers the noninvertibility directly.

The conditioning of the reduced system can, at the cost of size, be improved by avoiding certain unstable inversions in the recursive part. The paper will report how different invertibility criteria influence the reduced size and speed of convergence for some large economic planning models.

Session 2A. 11:20 - 11:40 a.m., Monday

SPARSE MATRIX TEST PROBLEMS

Iain S. Duff
AERE Harwell
Oxon, OX11 ORA, ENGLAND

Roger G. Grimes
John G. Lewis
W. G. Poole, Jr.
Boeing Computer Service, Inc.
Tukwila, Washington 98188

In the spring of 1982 we appealed in the IMANA and SIGNUM newsletters for further test matrices to augment collections at Harwell and BCS.

In this poster session, we report on the results of these appeals. Additionally, we indicate how we have organized the current set of test matrices and how we have classified the test matrices in the collection. We have designed our data base for ease of distribution, to facilitate the addition of further examples, and to allow the same matrix to be included in more than one classification.

We also explain the mechanism for submitting test examples for inclusion in the collection and for obtaining copies of the test matrices.

Poster Session. Monday and Tuesday

THE MULTIFRONTAL SOLUTION OF
UNSYMMETRIC SETS OF LINEAR EQUATIONS

Iain S. Duff
John K. Reid*
AERE Harwell
Oxon, OX11 0RA, ENGLAND

We show that general sparse sets of linear equations whose pattern is symmetric (or nearly so) can be solved efficiently by a multifrontal technique. The main advantages are that the analysis time is small compared to the factorization time and analysis can be performed in a predictable amount of storage. Additionally, there is scope for extra performance during factorization and solution on a vector or parallel machine. We show performance figures for examples run on the IBM 3033 and CRAY-1 computers.

Session 9A. 7:30 - 7:50 p.m., Tuesday

A PARTITIONING APPROACH TO PROCESSING NETWORKS AND THE SOLUTION OF SPARSE SYSTEMS

Michael Engquist
University of Texas
Austin, Texas 78712

Ordinary networks arise in connection with the transportation of commodities. Processing networks extend ordinary networks by allowing proportionality constraints among certain flows.

When the primal simplex algorithm is applied to a minimum cost flow problem for processing networks, the basis can be partitioned into a set of arcs which form a so-called representative spanning tree and a set of non-tree arcs. This partitioning yields a working basis which has a lower dimension than that which would result if a (non-representative) spanning tree were used. An implementation of the primal simplex algorithm has been developed which maintains the basis in this partitioned form. Computational results will be presented which show that this implementation compares very favorably with MINOS on a certain class of processing networks.

Sparse systems of equations can also be modelled as processing networks. It will be shown that, in some cases, solving the system involving the working basis is easier than solving the original system.

Session 12B. 11:40 - 12:00 noon, Wednesday

MATRIX MODIFICATION AND PARTITIONING

Albert M. Feltman
Boeing Computer Services Company
Tukwila, Washington 98188

The contingency analysis problem involves solving a sequence of systems of sparse equations where the matrices differ by a matrix of low rank. These systems can be solved by partitioning the matrix and confining the changes to the border, or by using the matrix modification formula to obtain the solution of one problem in terms of the previous one.

Often the low rank change has a special form: the matrix can be reordered so that all changes are confined to the lower right corner block. If this reordering is actually done, very little work is required in obtaining the solution of the modified problem based on the solution of the original problem. Unfortunately, this places a severe constraint on the partitioned form which is greatly amplified when a sequence of such changes in different parts of the matrix are desired. We will show a way of solving this problem, taking advantage of the special form of the low rank change, but not constraining the ordering of the matrix.

Session 9A. 8:10 - 8:30 p.m., Tuesday

A STRUCTURALLY STABLE MODIFICATION OF THE HELLERMAN-RARICK ALGORITHM FOR REORDERING UNSYMMETRIC SPARSE MATRICES

Albert M. Erisman
Roger G. Grimes
John G. Lewis*
William G. Poole, Jr.
Boeing Computer Services Company
Tukwila, Washington 98148

The P^u algorithm of Hellerman and Rarick has been used to reorder unsymmetric sparse matrices in order to minimize computation and storage costs when solving sparse systems of linear equations. Unfortunately, it is not infallible. It is known that the algorithm can generate intermediate matrices which are structurally singular and which may lead to a breakdown in the elimination process.

In this paper we present the algorithm in a top-down, block form and explain several of the problems which may occur. We describe several past attempts which were unsuccessful in correcting the problems.

We then define a new modification for handling the difficulties. This revised version of the algorithm will never produce structurally singular intermediate matrices if the original matrix is not structurally singular. Test results comparing this version with others will be given.

Session 2A. 11:00 - 11:20 a.m., Monday

SOLVING SPARSE STAIRCASE SYSTEMS BY GAUSSIAN ELIMINATION

Robert Fourer
Northwestern University
Evanston, Illinois 60201

A system of linear equations has a "sparse staircase" structure if its variables fall into a certain sequence of stages, such that any equation involves relatively few variables from at most two adjacent stages. Sparse staircase systems commonly arise in solving large multi-period linear programs by the simplex method.

This paper describes two approaches to solving sparse staircase systems by Gaussian elimination. Each approach combines techniques for stable period-by-period elimination of arbitrary staircase matrices with techniques for efficient elimination of arbitrary sparse matrices. The two approaches differ in their choice of sparse-elimination techniques: one uses "merit" techniques that select an attractive pivot at each elimination step, while the other employs "bump-and-spike" techniques that permute the staircase matrix to an attractive form in advance of elimination.

Sparse staircase elimination methods admit especially efficient implementations that operate on only a period or two of the staircase at a time. Their period-by-period elimination ordering can also be advantageous in handling certain sparse right-hand-side and solution vectors, and in computing solutions for subsets of periods. They may also produce an especially sparse factorization of the staircase matrix, although in general they do not find a sparser factorization than general sparse-elimination methods.

Session 2A. 11:40 - 12:00 noon, Monday

AN INCOMPLETE FACTORIZATION METHOD FOR SINGULAR IRREDUCIBLE M-MATRICES

R. E. Funderlic
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

R. J. Plemmons
North Carolina State University
Raleigh, North Carolina 27650

Markov queueing networks can give rise to very large, sparse, irreducible, singular (zero column sums) M-matrices A . The stationary probability distribution vector is the solution to a homogeneous system of linear equations $Ax = 0$. Certain economic petroleum reservoir and discrete Neumann problems give rise to related systems. The splitting $A = M - N$ with a matrix M having symmetric zero structure can be shown to be a regular splitting. A sparse LU factorization of a symmetric permutation, PM^T , is obtained using a standard symmetric ordering such as minimum degree. No pivoting for stability is necessary. Splitting strategies including those that have the larger $|a_{ij}|$ in M will be discussed.

Poster Session. Monday and Tuesday

ROW ORDERING SCHEMES FOR SPARSE GIVENS TRANSFORMATIONS

Alan George
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

Joseph Liu*
York University
Downsview, Ontario, Canada M3J 1P3

Let A be an m by n matrix, $m > n$, and let P_1 and P_2 be permutation matrices of order m and n respectively. Suppose $P_1 A P_2$ is reduced to upper trapezoidal form $\begin{pmatrix} R \\ 0 \end{pmatrix}$ using Givens rotations. The sparsity structure of R depends only on P_2 . For a given P_2 , the number of arithmetic operations required to compute R depends on P_1 . In this paper we consider row ordering strategies that are appropriate when P_2 is obtained from nested dissection type orderings of $A^T A$. Recently it was shown that the so-called "width-2" nested dissection orderings of $A^T A$ could be used to simultaneously obtain good row and column orderings for A . In this paper we show that the conventional (width-1) nested dissection orderings can also be used to induce good row orderings. Our analysis employs a bi-partite graph model of Givens rotations applied to A , similar in some respects to the bi-partite graph models of Gaussian elimination developed by Gilbert and Maulino.

Session 8B. 12:10 - 12:30 p.m., Tuesday

SOLUTION OF SPARSE UNDERDETERMINED SYSTEMS OF LINEAR EQUATIONS

Alan George
Esmond Ng*
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

In this paper we consider the problem of computing the minimal ℓ_2 -solution to a consistent underdetermined linear system $Ax = b$, where A is m by n with $m \leq n$. The method of solution is to reduce A to lower trapezoidal form $[L \ 0]$ using orthogonal transformations, where L is m by m and lower triangular. The method can be implemented efficiently if the matrix AAT is sparse. However, if A contains some dense columns, AAT may be unacceptably dense. We will present a method for handling these dense columns. The problem of solving a rank-deficient under-determined system will also be considered.

Session 8B. 12:30 - 12:50 p.m., Tuesday

NESTED DISSECTION WITH PARTIAL PIVOTING

John R. Gilbert*
Cornell University
Ithaca, New York 14853

Robert Schreiber
Stanford University
Stanford California 94305

When a sparse system of linear equations is solved by Gaussian elimination, many of the zero coefficients often become nonzero. One strategy for limiting this fill-in is nested dissection, which is applicable to many two-dimensional finite element and finite difference problems. Alan George invented this algorithm, and Lipton and Tarjan extended it to give fill-in that is within a constant factor of minimum for any system whose coefficient matrix represents a planar or near-planar graph.

Most analyses of nested dissection avoid numerical problems by assuming that the matrix of coefficients is symmetric and positive definite. What happens if the matrix is not so well behaved? This talk will present an algorithm that applies nested dissection while performing partial pivoting to control numerical stability. We shall analyze this "dissection pivoting" algorithm by using a bipartite graph model. For a large class of matrices, including those that represent planar graphs of bounded degree, the fill-in from this algorithm is within a constant factor of minimum. The constant is large, and we shall discuss what can be done to make the algorithm practical.

Session 9A. 7:50 - 8:10 p.m., Tuesday

IMPLEMENTATION OF A LARGE-SCALE LINEAR PROGRAMMING PACKAGE ON A MINICOMPUTER

James E. Giles
Tennessee Valley Authority
Norris, Tennessee 37828

Implementation of the linear programming portion of the software package MINOS on the HP1000F minicomputer is discussed. Memory restrictions are circumvented using the extended memory addressing capabilities of the HP1000F. An application of the implemented package to a water resource management problem is described.

Poster Session. Monday and Tuesday

ANALYSIS OF A MULTIGRID METHOD AS AN ITERATIVE TECHNIQUE
FOR SOLVING LINEAR SYSTEMS

Anne Greenbaum
Lawrence Livermore National Laboratory
Livermore, California 94550

A general class of iterative methods is introduced for solving symmetric, positive definite linear systems. These methods use two different approximations to the inverse of the matrix of the problem, one of which involves the inverse of a smaller matrix. It is shown that the methods of this class reduce the error by a constant factor at each step and that under "ideal" circumstances this constant is equal to $(\kappa' - 1)/(\kappa' + 1)$, where κ' is the ratio of the largest eigenvalue to the $(J+1)^{\text{st}}$ eigenvalue of the matrix, J being the dimension of the smaller matrix involved. A multigrid method is presented as an example of a method of this class, and it is shown that while the multigrid method does not quite achieve this optimal rate of convergence, it does reduce the error at each step by a constant factor independent of the mesh spacing h .

Session 4A. 7:50 - 8:10 p.m., Monday

PROGRAM FOR THE COMPARISON OF REORDERING ALGORITHMS FOR THE
SOLUTION OF UNSYMMETRIC SPARSE SYSTEMS OF EQUATIONS

Roger G. Grimes
John G. Lewis
William G. Poole, Jr.
Boeing Computer Services Company
Tukwila, Washington 98188

The authors have developed a program to monitor the effectiveness of various reordering algorithms for the solution of unsymmetric systems of equations over a broad collection of test problems. The objective is to select implementation-independent measures which can lead to the classification of matrices into categories where a particular reordering is most effective. Measures chosen include the amount of fill generated, operation counts, and numerical stability in the factorization of the permuted matrix.

A large collection of test matrices from diverse applications is being assembled to be used in analyzing matrix patterns. A description of the design and implementation of the program along with samples of preliminary results will be presented.

Poster Session. Monday and Tuesday

A NEW VERSION OF TARJAN'S STRONG CONNECT ALGORITHM

Fred G. Gustavson
IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

We describe a new version of Tarjan's algorithm for finding the strong components of a directed graph G with n nodes and N edges. The new version, called BLTF (Block Lower Triangular Form), is presented in a structured high level form that is translated into an efficient FORTRAN code. This version runs faster than a previous version called STCO. This algorithm is closely related to the algorithm MC13D of Duff and Reid, which is also an implementation of STRONG CONNECT. Algorithm BLTF is superior to MC13D in the following ways: it executes faster and uses less storage; it computes exactly the same output as Tarjan's algorithm; its translation into FORTRAN is completely structured; its inner loop is more efficient in that it asks half as many questions as MC13D as $N \rightarrow n^2$.

Session 7A. 11:45 - 12:00 noon, Tuesday

6.6

AN ALGORITHM FOR COMPUTING A FULL ASSIGNMENT FOR A SPARSE MATRIX

Fred G. Gustavson
IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

We describe an efficient implementation of the Assign Row Algorithm of Gustavson. This algorithm finds a maximal assignment for an arbitrary sparse (0-1) matrix and is based on the algorithm of M. Hall. This algorithm is closely related to algorithm MC21A of Duff, which is also an implementation of Hall's algorithm. Two versions of Assign Row are discussed. The first computes all cheap assignments before starting any depth first searches. The second implements Assign Row. Both algorithms are completely structured and efficiently translated into FORTRAN. The first algorithm almost always computes faster than the second and both algorithms almost always computes faster than MC21A.

Poster Session. Monday and Tuesday

AN IMPROVED ASSIGNMENT ALGORITHM

Fred G. Gustavson*
Gary D. Hachtel
IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

Aram K. Kevorkian
General Atomic Company
San Diego, California 92138

The Assign Row Algorithm of Gustavson uses the depth-first search technique to stretch an assignment in an undirected bipartite graph. An efficient implementation of this algorithm is the fastest method for finding an assignment for almost all graphs.

We present a new and improved version of the Assign Row Algorithm which performs equally well. However, on the small subset of all graphs where the Assign Row Algorithm performs poorly, the new algorithm runs up to an order of magnitude faster. The main feature of the new algorithm is the introduction of the efficient methods in the depth-first search procedure so that negative results in earlier depth-first searches can be used to reduce present and future depth-first searches. The new algorithm sheds insight into the problem of finding the block lower triangular form of a sparse matrix as follows. A partial assignment gives rise to a 'partial directed' graph and depth-first search can be used to quickly explore this graph. Failure to stretch an assignment in this graph leads to partial information on the block lower triangular form. Hence these parts of the graph need not be explored in later depth-first searches. Furthermore, the basic technique used in the two stage procedure of finding an assignment followed by the strong components determination of the associated directed graph merges into one idea of repeated depth-first search on a partial directed graph.

Session 7A. 11:00 - 11:20 a.m., Tuesday

A 1-PASS PROCEDURE FOR MAXIMAL ASSIGNMENT AND FINDING THE BLT FORM

Gary D. Hachtel*
Fred G. Gustavson
IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

Aran K. Kevorkian
General Atomic Company
San Diego, California 92134

A procedure is proposed for obtaining both a maximal assignment and the canonical block triangular form of a sparse matrix. In contrast to the conventional 2-pass approach, the proposed unified procedure incorporates block identification tests into the assignment algorithm. Thus blocks of the BLT form are discovered while in the process of solving the $O(|V|^2(|V|+|E|))$ maximal assignment problem.

This procedure is shown to simplify and decompose the maximal assignment problem by virtue of removing large row and columns blocks from further consideration. Each irreducible block is identified once only. The accumulated cost of the test is linear (i.e., $O(|V|+|E|)$). The new procedure is based upon DFS of the corresponding bipartite graph. Alternating edges in the palm tree thus obtained give a partial assignment of greater cardinality than that obtained by previously published "cheap assign" methods. Since the running time of maximal assignment algorithms depends chiefly on the number of subsequent "stretching" DFS passes, this reduces the expense of the overall procedure.

The partial assignment makes the bipartite graph partially directed. When backtracking during DFS of the bipartite graph, we perform tests analogous to Tarjan's transitive tests for biconnected components of the bipartite graph and strong components of the directed subgraph (both row-wise and col-wise). These tests can identify irreducible blocks of the BLT form (or sets of such blocks) prior to completing the assignment. We also perform certain "pre-minimum cover" tests which can identify such blocks without the requirement of backtracking. These tests require less computational expense, since they are not transitive.

We present some statistical data which suggest that the procedure might be favorable when applied to large problems with significant irreducibility.

A SPARSE LINEAR PROGRAMMING SUBPROGRAM

R. J. Hanson
K. L. Hiebert*
Sandia National Laboratories
Albuquerque, New Mexico 87185

We will describe a subprogram, SPLP, for solving linear programming problems. The package of subprogram units comprising SPLP is written in FORTRAN 77.

The subprogram SPLP is intended for problems involving at most a few thousand constraints and variables. The subprograms are written to take advantage of sparsity in the constraint matrix.

A very general problem statement is accepted by SPLP. It allows upper, lower, or no bounds on both the variables and the constraints. Both the primal and dual solutions are returned as output parameters. The package has many optional features. Among them is the ability to save partial results and then use them to continue the computation at a later time.

Session 12B. 10:40 - 11:00 a.m., Wednesday

SPARSE ORTHOGONAL SCHEMES FOR STRUCTURAL OPTIMIZATION USING THE FORCE METHOD

M. L. Miller
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

R. J. Williams*
North Carolina State University
Raleigh, North Carolina 27651

R. C. Ward
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

The force method enables one to calculate the system force vector F acting on a structure, and the associated stresses and strains, without reassembling and refactoring the overall stiffness matrix at each step of the optimization process. The method consists of two parts: (1) computation of a matrix B whose columns form a basis of the null-space of the equilibrium matrix E , which is $m \times n$ of rank m , along with a particular solution S to the underdetermined system $ES = P$, where P is the set of external loads on the structure, and (2) the solution of the least squares problem $\min_x \|f(Bx + S)\|_2$, where

$f = \tilde{f}^T \tilde{f}$ is the element flexibility matrix which changes at each optimization step. In this case $F = Bx + S$ satisfies the principle of complementary potential energy - that $E^T F$ is minimal over all F such that $EF = P$.

In this talk, an orthogonal decomposition scheme for solving part (1), called the turnback-QR method, which preserves the banded structure of E in computing B , is presented. Some preliminary comparisons are made with the turnback-LL method based upon elimination, given recently by Kaneko, et al [1982]. Also, some comments are made on the use of the linearly constrained suboptimal gradient method of Powell for computing x in (2) at each optimization step.

Session 3B. 12:30 - 12:50 p.m., Monday

INCOMPLETE METHODS FOR SOLVING $A^T A x = b$

J. D. Evans¹
Queens University,
Belfast, Northern Ireland

M. A. Ajaz
University of Engineering and Technology
Lahore, Pakistan

Much interest has recently been shown in incomplete factorization methods for solving sparse linear simultaneous equations. They have the advantage of avoiding most or all of the fill-in associated with pure elimination methods and, although iterative in nature, they exhibit much better convergence rates than do classical iterative methods.

In this paper incomplete methods are considered for the solution of $A^T A x = b$ where A is sparse. One possibility is to construct the product of $A^T A$ and then use an ICCG algorithm for their solution. Alternatively it is possible to perform incomplete orthogonal decompositions of A by modifying either the Gram-Schmidt or the Givens rotation methods. These decompositions will yield incomplete triangular factors for $A^T A$ which differ from the one obtained by ICCG, but which give a similar iterative phase for the solution process.

Implementations of each of these three methods are briefly described and some numerical comparisons made using examples from structural analysis. The ICCG method has the advantage that it is easy to implement using packing schemes for the elements of the sparse matrices. The orthogonal decomposition methods, particularly the one using Givens rotations, do appear to have somewhat better numerical efficiencies when counting only non-zero arithmetical operations, but it is not obvious how their decomposition phases could be efficiently implemented in a sparse store.

Session 12A. 11:00 - 11:20 a.m., Wednesday

A PRODUCTION EIGENANALYSIS SYSTEM FOR LARGE, GENERALIZED SYMMETRIC PROBLEMS

Paul S. Jensen
Lockheed Palo Alto Research Laboratory
Palo Alto, California 94304

A system of programs called BES (Basic Eigenanalysis System) for determining a few eigenpairs of problems of the form

$$Ax = Mx \lambda \quad (1)$$

has been developed for production engineering analysis. A and M are typically large, real, symmetric and sparse matrices for which there exist scalars α and μ such that $\alpha A - \mu M$ is positive definite. The eigenvalues of such problems are real and the desired eigenpairs are typically specified by giving the limits of a desired spectral range (or section).

The system is organized as independent executable programs that operate in a chained (pipeline) fashion, communicating via a database system. For each problem analyzed, the chaining details are developed dynamically by the system to suit the particular problem and the results are deposited in a global database.

For the solution, a series of problems of the form

$$(A - \sigma M)y = (\alpha A - \mu M)y \quad (2)$$

is solved, where α and μ are selected once for positive definiteness and "shift point" σ varies. Each problem in form (2) is transformed to a classical form and solved using a block Lanczos algorithm, originated by Parlett and Scott, that is specifically designed for spectral sections. The transformations require one factorization of $\alpha A - \mu M$ and one factorization of $A - \sigma M$ for each value of σ .

Each shift point σ is at the center of a subsection (interval) within the range specified by a user. For each subsection, the Lanczos analysis either determines all of the eigenvalues contained therein or it determines several in the vicinity of the shift point. In the former case, no further analysis of the subsection is performed. In the latter case, two new subsections (at the lower and upper extremes of the given subsection) are generated and subsequently analyzed. The usual checks for completeness of coverage are utilized and refinements via Ritz projections are made when needed.

A variety of problems from structural vibration and buckling analysis have been solved using BES. It appears to perform typical analyses for about one-sixth of the cost attributed to previous programs based on simultaneous inverse iteration.

THE CLASSES $M_{n,k}$ AND PROPERTIES OF SPARSE MATRICES
RESEMBLING THOSE OF MATRICES IN A SMALLER DIMENSION

C. R. Johnson
University of Maryland
College Park, Maryland 20742

For a 2-by-2 component-wise nonnegative matrix, both eigenvalues are real and the spectral radius is easily calculated or estimated. For an n -by- n tri-diagonal nonnegative matrix these observations generalize: it is well known that the eigenvalues are all real, and it may be shown that the spectral radius is on the order of (in fact, between 1 and 2 times) the largest of the spectral radii of the 2-by-2 principal submatrices occurring consecutively along the diagonal. These facts support the, perhaps often made, meta-observation that properties of tri-diagonal matrices are very much like those of 2-by-2 matrices. However, both these facts generalize considerably further in addressing the issue of what it is about the sparsity pattern of a large matrix with many zeroes which make it "like" a matrix in a much smaller dimension. If the longest simple circuit in the usual directed graph of an n -by- n matrix is k ($\leq n$), we shall say that it belongs to the class $M_{n,k}$. For example $M_{n,0}$ are the nilpotent essentially triangular matrices $M_{n,1}$ are the remaining essentially triangular matrices; and $M_{n,2}$ includes (but is not restricted to) the nontrivial tri-diagonal matrices. We present here some recent results about ways in which $M_{n,k}$ is more like the k -by- k matrices than the n -by- n matrices, regardless of the size of n . These include spectral properties of nonnegative matrices and general matrices, D -stability, determinantal calculations, and others. Often, attempts are made to generalize tri-diagonal results to penta-diagonal matrices, etc., just as 2-by-2 results might be generalized to the 3-by-3 case. However, if the tri-diagonal fact is actually an $M_{n,2}$ fact, generalization to $M_{n,r}$ (the penta-diagonals are generally in $M_{r,n}$) may be more fruitful.

Session 8A. 12:50 - 1:10 p.m., Tuesday

AN EFFICIENT ALGORITHM FOR THE FACTORIZATION OF POSSIBLY RANK DEFICIENT MATRICES

M. S. Kamel*

K. Singhal

University of Waterloo
Waterloo, Ontario, Canada N2L 3G1

Rank deficient matrices arise due to the problem's inherent structure, or due to data uncertainties or perturbations induced by round-off errors. It is desired to have algorithms which can detect the possibility of rank deficiency of a matrix in order to replace it by an actual rank deficient matrix. Lawson and Hanson propose an algorithm that uses Householder transformations with column interchange to triangularize the matrix and determine its pseudorank. Their algorithm is mainly focused on solving least squares problems.

In this paper we propose a similar algorithm to factorize the matrix using Householder transformations also with column interchanges. However, the new algorithm avoids operations on those parts of the matrix that will eventually be discarded due to rank deficiency. The algorithm offers computational savings in the general case, but its major advantage is that it can be applied to sparse matrices without the risk of ruining the sparsity pattern that is usually experienced when applying Householder transformations to sparse matrices.

In this paper we present the algorithm, operation count analysis, applications to least squares problems, and other applications. Numerical experiments and examples will be reported.

Session 8B. 12:50 - 1:10 p.m., Tuesday

75

A THREE DIMENSIONAL FINITE ELEMENT MAGNETOTELLURIC MODELING PROGRAM

W. P. Kamp
Phillips Petroleum Company
Bartlesville, Oklahoma 74004

A three-dimensional modeling program for magnetotellurics has been written using the finite element method. The program has been successfully implemented on the CRAY and IBM 3033 computers. We will discuss the experience gained by this exercise both from a numerical analysis and geophysical point of view. Example models together with comparisons to known models will be presented.

The use of various sparse matrix packages in solving this problem will be discussed. Of particular interest will be the application of complex versions of the algorithms on the CRAY computer.

Session 9B. 8:10 - 8:30 p.m., Tuesday

THE TURN-BACK LU PROCEDURE FOR COMPUTING A SPARSE AND BANDED BASIS OF THE NULL SPACE

Ken Kaneko
University of Wisconsin
Madison, Wisconsin 53706

This talk is concerned with the problem of computing a basis of the null space of a given large sparse and banded matrix A with full row rank. The problem is not only one of the basic problems in linear numerical analysis, but also has many applications of practical importance. Two significant such applications are (i) structural analysis, and (ii) a new implementation of the Ellipsoid Algorithm for solving a linear system, or a pair of dual linear programs.

Determining any basis of the null space of a given matrix A is not a particularly difficult task and there are many existing methods to do the task. Existing methods, however, tend to destroy either sparsity and/or the banded nonzero pattern of A . In a recent paper, we have shown that a method which we call the Turn-Back LU Procedure is extremely effective in computing a basis of the null space of A , while retaining the sparsity and bandedness. We shall show how this Turn-Back LU Procedure works and demonstrate its superiority over a few existing methods by means of several examples arising from structural analysis.

Session 3B. 12:10 - 12:30 p.m., Monday

FAST DIRECT METHODS FOR
OTHER SPARSE MATRIX PROBLEMS

Linda Kaufman
Bell Laboratories
Murray Hill, New Jersey 07974

An $n \times n$ matrix A is separable if it can be written as $A = \sum_{i=1}^m B_i \otimes C_i$

where the dimension of the B_i 's is greater than 1 and there exist matrices Q and Z such that for all i , QB_iZ are diagonal matrices. Fast direct methods based on matrix separability have been used to solve the linear system arising for a finite difference discretization of Poisson's equation. We will show how matrix separability can be used to formulate efficient algorithms in two other contexts. The first system arises when applying Galerkin's method to solve separable elliptic p.d.e.'s on a rectangle while using a tensor product B-spline basis. Using the separability of the system produces a dramatic decrease in time and space requirements over traditional sparse direct methods. The second situation arises when determining the probability distribution during a queuing analysis of a network. Because many of the diagonal submatrices of the matrix defined by the Kolmogorov balance equations are often separable, a fast block iteration technique can be formulated.

Session 3A. 12:50 - 1:10 p.m., Monday

A NEW THEORY ON PERMUTING MATRICES TO BLOCK TRIANGULAR FORM

Armen K. Keownian*

General Atomic Company
San Diego, California 92138

Fred G. Gustavson

Gary D. Hachtel

IBM T. J. Watson Research Center
Yorktown Heights, New York 10598

In this paper we introduce the concept of block transversal in a matrix. We use this concept to establish a series of results on block triangular matrices with fully irreducible (indecomposable) diagonal blocks. Our results include a new theorem on the uniqueness of the block triangular form. Additionally, we employ the concept of block transversal to generalize a generally accepted two-step procedure for computing the block triangular form of a square matrix.

Session 7A. 10:40 - 11:00 a.m., Tuesday

BLOCK ITERATIVE METHODS WITH AGGREGATION FOR SOLVING NEARLY COMPLETELY DECOMPOSABLE MARKOV CHAINS

Robert Koury
David F. McAllister
William J. Stewart*
North Carolina State University
Raleigh, North Carolina 27650

Iterative methods have long been used to obtain the stationary probability vector of Markov chains. These chains give rise to stochastic matrices which are often very large and extremely sparse. Furthermore, in most practical applications the matrices possess a distinctive non-zero structure. The Markov chain is said to be nearly completely decomposable when it is possible to symmetrically permute the matrix so that the probability mass is concentrated into diagonal blocks; the non-zero elements of the off-diagonal blocks being relatively small in magnitude. Under these circumstances the rate of convergence of the usual iterative methods is so slow that they are of no practical value and analysts have been obliged to turn to decompositional and aggregative techniques to determine approximate solutions. The accuracy of the approximation obtained is in general proportional to the degree to which the probability mass is concentrated into the diagonal blocks.

Recent research has lead to a number of methods in which aggregation is incorporated into iterative procedures to enable exact solutions to be computed efficiently. In this paper we distinguish two different aggregation techniques and discuss their applicability. We show how these techniques may be combined with the group iteration methods of Gauss-Seidel and successive overrelaxation, and we present the results which were obtained from a number of important examples.

Session 3A. 12:10 - 12:30 p.m., Monday

A COMPARISON OF TWO SPARSE MATRIX
PROCESSING TECHNIQUES FOR
STRUCTURAL ANALYSIS APPLICATIONS

William A. Loden
Lockheed Missiles and Space Company, Inc.
Sunnyvale, California 94086

The performance of algorithms based on two techniques for the solution of large sparse, symmetric linear equation systems of the type that arise frequently in structural analysis are studied. Representative test problems, for simple and complex structural configurations treated with the finite element method, are solved with two implementations of Cholesky method profile algorithms and with the SPSYST sparse-matrix system package developed by J. K. Reid.

Poster Session. Monday and Tuesday

CONJUGATE GRADIENT IMPLIES NORMAL

Tom Manteuffel*
Vance Faber
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

A conjugate gradient-like iteration for general linear systems has long been sought. The existence of such a method depends upon how "conjugate gradient-like" is defined. One can always form the normal equations and apply the conjugate gradient method for symmetric positive definite systems. We consider gradient methods as defined by Rutishauser; that is, each iterate must come from the appropriate Krylov subspace associated with the matrix A . Further, we insist that each iterate be optimal over the Krylov subspace with respect to a norm associated with the inner product. (The inner product may depend upon A .) Under these assumptions it can be shown that an s -term conjugate gradient method exists for every initial error if, and only if,

- 1) the minimal polynomial of A is of degree $\leq s$
- 2) A is normal with respect to the inner product and A^* is a polynomial in A of degree $\leq s-2$ in any orthogonal basis imposed by that inner product.

This implies that a 3-term conjugate gradient iteration exists if, and only if, $A^* = A$ or $A = dI - B$, $B^* = -B$, or the degree of the minimal polynomial of A is ≤ 3 .

This talk will sketch a proof of these results and discuss implications and extensions.

Session 4A. 7:30 - 7:50 p.m., Monday

A FAST ALGORITHM
THAT MAKES MATRICES OPTIMALLY SPARSE

S. Thomas McCormick
Stanford University
Stanford, California 94305

Under a non-degeneracy assumption on the non-zero entries of a given sparse matrix, a polynomially-bounded algorithm is presented that performs row operations on the given matrix which reduce it to the sparsest possible matrix with the same row space. For each row of the matrix, the algorithm performs a maximum cardinality matching in the bipartite graph associated with a submatrix which is induced by that row. The dual of the optimal matching then specifies the row operations that will be done on that row. A variant of the algorithm that processes the matrix in place is also described. A particularly promising application of this algorithm is to reduce linear constraint matrices as a way of accelerating optimization.

Session 12B. 11:20 - 11:40 a.m., Wednesday

NULL-SPACE METHODS FOR LARGE-SCALE QUADRATIC PROGRAMMING

Walter Murray
Stanford University
Stanford, California 94305

Null-space methods form a powerful class of techniques for both convex and non-convex quadratic programming. We shall consider such methods for the solution of two classes of large sparse QP. The first class contains problems for which t , the number of constraints active at the solution, is close to n , the number of variables. The second class contains problems for which t is small relative to n . A dominant feature of the storage and computational overhead required for a null-space method is the solution of an $(n-t) \times (n-t)$ system of linear equations. It will be shown how such systems may be solved efficiently when t is small.

Session 12B. 11:00 - 11:20 a.m., Wednesday

A PORTABLE SOFTWARE PACKAGE FOR
NONLINEAR PARTIAL DIFFERENTIAL EQUATIONS

J. W. Neuberger
North Texas State University
Denton, Texas 76203

R. J. Renka*
Union Carbide Corporation, Nuclear Division
Oak Ridge, Tennessee 37830

This paper describes a storage-efficient method and associated software package for the solution of a general class of second-order nonlinear partial differential equations. The method is essentially an iterative scheme for the least squares minimization of a finite difference discretization of the residual. The advantages of this approach are its ability to treat a wide range of problems (it is type-independent) and its low storage requirements - $O(N)$ for N grid points. The effectiveness of the method is based on a gradient-smoothing technique which increases the rate of convergence of the iterative procedure.

Session 2B. 11:40 - 12:00 noon, Monday

ELEMENT PRECONDITIONING

Bahram Nour-Omid*
Beresford N. Parlett
University of California
Berkeley, California 94720

Elliptic boundary value problems are turned into large symmetric systems of equations. These large matrices are usually assembled from small ones. It is simple to omit the assembly process and use the code to accumulate the product Kv for any v . Consequently the conjugate gradient algorithm (CG) can be used to solve $Ku = f$ without ever forming K .

It is well known however that CG should be applied to preconditioned systems. In this paper we show how to achieve preconditioning without forming any large matrices.

The trade off between time and storage is examined for the 2-D model problem and the analysis of several realistic structures.

Session 12A. 11:40 - 12:00 noon, Wednesday

BANDED PRECONDITIONING FOR SOLUTION OF FINITE ELEMENT EQUATIONS

Bahram Nour-Omid*
Horst D. Simon
University of California
Berkeley, California 94720

The preconditioned conjugate gradient algorithm has been successfully applied to solving symmetric linear systems of equations arising from finite difference and finite element discretizations of a variety of problems.

In this paper we consider a matrix splitting, $A = M - R$, where M is the dense banded part of A . We identify a certain class of symmetric positive definite matrices for which M is also positive definite. This class contains matrices arising from finite element discretization of elliptic boundary value problems.

M was used as a preconditioning matrix for the conjugate gradient and the Lanczos algorithms to solve a variety of practical problems in structural engineering. Favorable results were obtained.

Session 12A. 10:40 - 11:00 a.m., Wednesday

SPARSE MATRICES IN SIMPLIFICATION OF SYSTEMS

A. S. Rao
University of Wisconsin - Parkside
Kenosha, Wisconsin 53140

In recent years simplification of higher order systems has been investigated by several authors and the simplified models are used for various design purposes. The controller design using simplified models for computer control application requires that the simplification algorithm employed to derive lower order models use minimal computer time and memory. It is known that the computational requirements of the Routh approximant simplification procedure are minimal.

To derive lower order models for stable systems the Routh approximant method involves transformation of the system description into γ - ξ canonic form and then suppression of less dominant modes. The simplification of unstable systems using Routh approximant procedure involves a translation equivalent to the shift of imaginary axis in the s plane, followed by a transformation of the modified system into γ - ξ canonic form.

In both cases the major computational requirements are involved in the evaluation and inversion of the appropriate transformation matrices, which are sparse and have special structure. The paper presents the details of these sparse matrices and a procedure to determine their inverses.

Session 3B. 12:50 - 1:10 p.m., Monday

PRECONDITIONING STRATEGIES FOR THE
CONJUGATE GRADIENT ALGORITHM ON MULTIPROCESSORS

A. Sameh*

C. Taft

University of Illinois
Urbana, Illinois 61801-2987

In this paper we consider the conjugate gradient (C.G.) algorithm, with preconditioning, for solving positive-definite linear systems on multiprocessors. We deal primarily with those systems that arise from the finite-difference discretization of self-adjoint elliptic boundary value problems in two and three dimensions. Assuming that the multiprocessor consists of a set of linearly connected processors, we investigate the organization of the C.G. algorithm so as to achieve maximum speedup over the sequential scheme. Preconditioning strategies that enhance the convergence of the C.G. algorithm are also investigated. For example, in the case of two-dimensional problems, preconditioning strategies that are suitable for our multiprocessor include: (a) block-Jacobi splitting in conjunction with line red-black ordering, (b) incomplete Cholesky factorization (for M-matrices) in conjunction with point red-black ordering, and (c) a block generalization of the incomplete Cholesky factorization in conjunction with two (or more) line red-black ordering. Furthermore, we explore the suitability of the above preconditioned C.G. schemes when the processors are not tightly coupled, i.e., when the transfer of one floating-point number from one processor to another is more costly than one arithmetic operation.

Session 4B. 8:10 - 8:30 p.m., Monday

LAS02--SPARSE SYMMETRIC EIGENVALUE PACKAGE

David S. Scott
University of Texas
Austin, Texas 78712

The LAS0 package (the Lanczos Algorithm with Selective Orthogonalization) became available through the National Energy Software Center at the Argonne National Laboratory in 1981. LAS0 used EISPACK subroutines to compute a few eigenvalues of a symmetric band matrix at each step of the algorithm. It soon became apparent that the EISPACK subroutines were not very efficient when the band matrix got very long compared to the width of the band. A new eigenvalue solver was developed based on the Rayleigh quotient iteration which overcame this difficulty. This paper describes the eigenvalue solver, its incorporation into the package and the other modifications which were made to the package at the same time.

Session 7B. 10:40 - 11:00 a.m., Tuesday

IMSL SOFTWARE FOR DIFFERENTIAL EQUATIONS IN ONE SPACE VARIABLE

Granville Sewell
IMSL, Inc.
Houston, Texas 77036

Two routines developed by the author for the IMSL Library are discussed.

DTPTB (Edition 8) solves boundary value problems using a multiple shooting technique, with IMSL routine DVERK employed to solve all initial value problems. When Newton's method is used to solve the resulting system of simultaneous equations, a Jacobian with a "periodic band" structure arises.

DPDES (Edition 9) solves a partial differential equation system of the form $u_t = f(x, t, u, u_x, u_{xx})$ using a collocation method and the method of lines. Cubic Hermite basis functions are used to integrate the method of lines ODE systems $y' = A^{-1}f(t, y)$, where although both A and the Jacobian of f are banded, A^{-1} is full. A very simple modification of DGEAR, which could be applied to other GEAR codes designed only for the usual ODE problem $y' = f(t, y)$, was necessary to handle this system without the need to deal with any full matrices.

Session 2B. 11:20 - 11:40 a.m., Monday

THE LANCZOS ALGORITHM WITH PARTIAL REORTHOGONALIZATION
FOR THE SOLUTION OF NONSYMMETRIC LINEAR SYSTEMS

Horst D. Simon
SUNY at Stony Brook
Stony Brook, New York 11794

The Lanczos algorithm with partial reorthogonalization has recently been applied by the author to the solution of large sparse symmetric linear systems of equations. Here this work is extended for the nonsymmetric case and linear least squares problems $Bx = f$. This is done by applying the symmetric algorithm to the problem

$$\begin{bmatrix} 0 & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} y \\ x \end{bmatrix} = \begin{bmatrix} f \\ 0 \end{bmatrix} .$$

as Paige [TOMS 8, Vol. 1, 1982] already showed, several simplifications result because of the special structure of the matrix and the right hand side. By using partial reorthogonalization a new algorithm is obtained which minimizes the number of matrix-vector multiplications. We present some numerical examples, which show that the algorithm works well.

Session 7B. 11:40 - 12:00 noon, Tuesday

SOME OPERATIONS ON SPARSE MATRICES

Daniel B. Szyld*
Oleg Vishnepolsky
Institute for Economic Analysis
New York University
New York, New York 10003

In many applications, the matrices are so large that even with a sparse matrix storage scheme, not more than one or sometimes two of them will fit in core. We developed an algorithm for permuted transposition when only the original matrix is kept in core and the resulting one is written on secondary storage. This algorithm is based on the distribution of all non-zero elements among the column queues, using a link vector.

A second algorithm, to perform the multiplication of two sparse matrices, overwrites the portion of the factors that has already been used, and if this amount of storage is not sufficient, it writes the product out of core. These two algorithms extend those presented by F. Gustavson [TOMS, 4, 1978, pp. 250-268]. We present some comparisons.

To obtain rows and/or columns permutations of a sparse matrix, one can execute the permuted transposition algorithm twice. We offer an alternative method that is faster if the resulting matrix is used with its indices unordered, and we explore different sorting techniques for ordering the indices. The same algorithm is used to obtain a submatrix (any subset of rows and columns) of a sparse matrix.

Session 13A. 12:30 - 12:50 p.m., Wednesday

INVERSE MATRIX REPRESENTATION WITH ONE TRIANGULAR ARRAY (IMPLICIT GAUSS)

Silvio Ursic
University of Wisconsin
Madison, Wisconsin 53705

This talk will describe a technique that permits the representation of the inverse of a matrix A with only one triangular array. More precisely, let $LA = U$ with L lower triangular and U upper triangular arrays. Then, to compute $A^{-1}b$ and b^TA^{-1} in $O(N^2)$ operations, we need the arrays:

- (1) L, U (Gauss);
- (2) L^{-1}, U (Crout);
- (3) L, A (Implicit Gauss).

The array L can be obtained, in $O(N^3)$ operations, without having to compute all the elements of U .

We will discuss some of the sparsity implications of being able to concentrate on obtaining a sparse array L , not caring about the number of non-zeros in the array U .

Session 8A. 12:10 - 12:30 p.m., Tuesday

